arXiv:1404.4079v1 [math.OC] 15 Apr 2014

Reconstructing trajectories from the moments of occupation measures

Mathieu Claeys

***Abstract*— Moment optimization techniques have been recently proposed to solve globally various classes of optimal control problems. As those methods return truncated moment sequences of occupation measures, this paper explores a numeric method for reconstructing optimal trajectories and controls from this data. In fact, by approximating occupation measures by atomic measures on a given grid, the problem reduces to a finite-dimensional linear program. In contrast with earlier numerical methods, this linear program is guaranteed to be feasible, no tolerance needs to be specified, and its size can be properly controlled. When combined with local optimal control solvers, this yields a powerful and flexible numerical approach for tackling difficult control problems, as demonstrated by examples.**

Moments optimization techniques have emerged recently as a versatile tool for the global resolution of many non linear optimization problems, see for instance [12] and references therein for applications to polynomial optimization, optimal control, stochastic processes, and more. The general procedure is illustrated in Fig. 1. First of all, the problem of interest is lifted or relaxed as a Linear Program (LP) on measures. As such, the problem becomes convex, albeit on a vector space which is hardly tractable in the general case, besides brute-force discretization. However, when problem data is polynomial, measures can be manipulated by their moments, which leads to a well studied hierarchy of moment relaxations, whose cost converges asymptotically to that of the measure LP.

After numeric resolution of one of those relaxations, it is of obvious interest to assert termination of the hierarchy, as well as extract from moment data the solution of the measure LP, which in turn must be interpreted in terms of solutions of the original problem. This last step of the method is a typical example of an inverse problem.

For polynomial, finite-dimensional optimization as proposed in [10], the full approach is now well mastered. From the optimized moment data, one can expect generically at a finite relaxation order [14] to recover the support of the measure, which is located at the (generically unique) optimal point. Numerical tools allow for the resolution of such polynomial problems by the mere definition of the problem (see for instance [8]), for small size problems.

For Optimal Control Problems (OCP), on the other hand, most recent work has been focused on the convex lift and optimization step of the approach, for various classes of systems: from the case of bounded controls as presented in Lasserre et al. [11], the approach was subsequently extended to impulsive controls [3] or switched systems [6]. However, very little work has been done on the inverse problem, which precludes the same "black box" operations as for polynomial optimization. To the author's knowledge, only exploitation of the dual problem as hinted in [11] or the resolution of an inverse problem with very limited moment data [9] have been investigated. We present briefly those methods and compare them to the one presented on this paper in §IV.

Several obstacles related to the nature of solutions of OCPs explain such a gap. First of all, trajectories are inherently infinite-dimensional objects, such that one can only expect to have approximate moments to the measure LP as a result of any given *finite order* moment relaxation. In addition, even if those moments were exact, their finite number prevents the exact reconstruction of an infinite dimensional object. Therefore, only approximate optimal trajectory and control law can be recovered. This is especially true for the control, which is expected to be highly discontinuous, switching in between bang, path-constrained and singular arcs.

Due to the restrictions mentioned above, we propose a numeric approach to reconstruct trajectories. In much of the spirit of the moment approach itself, we propose a convex program for extracting trajectory/control way points. As such, the scheme is guaranteed to converge if an admissible solution exists, and users should not specify any starting point for the method. The method is based on the approach developed by Rubio (see [16] and the references therein), which uses a linear program to solve the measure LP directly. However, the key difference here is that in our approach, only an inverse problem must be solved, whereas the actual trajectory optimization is performed through semi-definite relaxations. The segregation between optimization and inverse problem allow the overcome several difficulties arising in those earlier works.

As a last step, the obtained way points are used as starting data of a direct optimal control routine. This allows for the reconstructing of a trajectory to any given precision. As a result, this also equips to approach with a termination criterion; At any given relaxation order, if the cost of the solution of the direct method agrees with some prescribed tolerance to the cost of the moment relaxation, the solution is certified as globally optimal.

The idea of using a global method for hot-starting local optimal control methods is surprisingly not well studied in the literature. We explain this with the fact that the numeric resolution of the Hamilton-Jacobi-Bellman (HJB) equation is an endeavor that can be hardly undertaken gradually in function of computational power available, as bounds on the discretization grid must be enforced to avoid numerical

M. Claeys is with Department of Engineering, University of Cambridge, Trumpington Street, Cambridge CB2 1PX, United Kingdom, mathieu.claeys@eng.cam.ac.uk

instability. In addition, once the solution is available, only those applications requiring high-precision solutions truly need a second round of computations from a local method. This is the motivation behind Cristiani and Martinon [4], who propose to solve the HJB equation on a coarse grid to initialize an indirect method exploiting Pontryagin's maximum principle [15], via their well-known relationship (see e.g. [2] in the constrained case).

In contrast with these limitations, moment relaxations of low order are very cheap to compute, and as shown in the examples, are usually tight enough to hot start a local method in a sufficiently close neighborhood of the optimal solution. Up to the author's knowledge, this is the first time that moment relaxations are used for such a purpose.

NOTATIONS

Let $\mathbf{Z} \in \mathbb{R}^n$ be a compact set of an Euclidean space. We note by $\mathcal{M}^+(\mathbf{Z})$ the space of finite, positive measures supported on $\mathbf{Z}$, equipped with the weak-$*$ topology. For a continuous function $f(z) \in C(\mathbf{Z})$, denote by $\int_{\mathbf{Z}} f(z)\,\mu(dz)$ the integral of $f(z)$ by the measure $\mu \in \mathcal{M}^+(\mathbf{Z})$. When no confusion may arise, we note $\langle f, \mu\rangle$ for the integral to simplify exposition and to insist on the duality relationship between continuous functions and measures. The Dirac measure supported at z^* is denoted by δ_{z^*}.

For multi-index $\alpha \in \mathbb{N}^n$ and vector $z \in \mathbb{R}^n$, we use the notation $z^\alpha := \prod_{i=1}^n z_i^{\alpha_i}$. Denote by $\mathbb{N}^n_m$ the set $\{\alpha \in \mathbb{N}^n : \sum_{i=1}^n \alpha_i \leq m\}$. The moment of multi-index $\alpha \in \mathbb{N}^n$ of measure $\mu \in \mathcal{M}^+(\mathbf{Z} \subset \mathbb{R}^n)$ is then defined as the real $y_\alpha = \langle z^\alpha, \mu\rangle$. A multi-indexed sequence of reals $\{y_\alpha\}_{\alpha\in\mathbb{N}^n}$ is said to have a *representing measure* on $\mathbf{Z}$ if there exists $\mu \in \mathcal{M}^+(\mathbf{Z})$ such that $y_\alpha = \langle z^\alpha, \mu\rangle$ for all $\alpha \in \mathbb{N}^n$.

Denote by $\mathbb{R}[z]$ the ring of polynomials in the variables z. A set $\mathbf{Z} \in \mathbb{R}^n$ is basic semi-algebraic if it is defined as the intersection of finitely many polynomial inequalities: $\mathbf{Z} := \{z \in \mathbb{R}^n :\ g_i(z) \geq 0,\, g_i(z) \in \mathbb{R}[z],\, i = 1 \dots n_{\mathbf{Z}}\}$.

Finally, we use the notation $\underline{x}$ to denote parameters of the state space where trajectories $x(t)$ live. We use the same convention $\underline{u}$ for the controls $u(t)$. This notation makes the passage from temporal integration to integration with respect to a measure transparent.

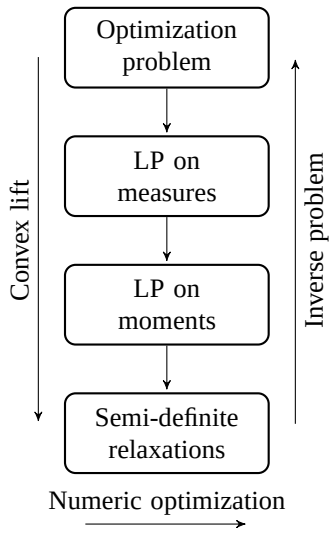

Fig. 1. The moment approach. An optimization problem is relaxed as semi-definite relaxations, solved numerically. The solution of the original problem is found by solving an inverse problem from the optimized data.

I. MOMENT OPTIMIZATION FOR OPTIMAL CONTROL

This section outlines the main steps of the convex lift step of the moment approach for bounded OCPs as presented in [11]. We only highlight the important features necessary for the following sections.

Consider the following end-constrained[1] problem:

$$\begin{aligned} J = \inf_{u(t)} & \int_{t_i}^{t_f} h(t,x,u)\,\mathrm{d}t \\ \text{s.t. } & \dot{x} = f(t,x,u), \\ & x(t_i),\ x(t_f) \text{ given}, \\ & x(t) \in \mathbf{X}, \quad u(t) \in \mathbf{U}, \\ & t \in \mathbf{T} := [t_i, t_f], \end{aligned} \tag{1}$$

where the state $x(t) \in \mathbb{R}^n$, controls $u(t) \in \mathbb{R}^m$ and functions $f, h \in \mathbb{R}[t, \underline{x}, \underline{u}]$. Sets $\mathbf{X}$ and $\mathbf{U}$ are compact, basic semialgebraic sets, and are assumed w.l.g. to be contained in a ball included in the algebraic definition of those sets.

For a process $(u(t), x(t))$ admissible for (1), the *occupation measure* $\mu \in \mathcal{M}^+([\mathbf{T} \times \mathbf{U} \times \mathbf{X})$ is defined by:

$$\mu(\mathbf{A} \times \mathbf{B} \times \mathbf{C}) := \int_{\mathbf{A}} \delta_{u(t)}(\mathbf{B})\,\delta_{x(t)}(\mathbf{C})\,\mathrm{d}t, \tag{2}$$

where $\mathbf{A}$, $\mathbf{B}$ and $\mathbf{C}$ are Borel subsets of resp. $\mathbf{T}$, $\mathbf{U}$ and $\mathbf{X}$.

When evaluating a continuously differentiable test function $v(t, \underline{x})$ along such an admissible trajectory, straightforward computations reveal that the occupation measure is admissible for the following measure LP, while achieving the same cost as (1):

$$\begin{aligned} J_{LP} = \inf_{\mu} & \langle h, \mu\rangle \\ \text{s.t. } \forall v \in \mathbb{R}[t,\underline{x}] :\ & [v(\cdot, x(\cdot))]_{t_i}^{t_f} = \langle \frac{\partial v}{\partial t} + \frac{\partial v}{\partial \underline{x}} f, \mu\rangle, \\ & \mu \in \mathcal{M}^+(\mathbf{T} \times \mathbf{U} \times \mathbf{X}). \end{aligned} \tag{3}$$

Obviously, since the admissible elements have been enlarged, $J_{LP} \leq J$. It is however expected generically that $J_{LP} = J$ hold, see e.g. the discussion in [7].

As problem data was assumed polynomial, viz. $f, h \in \mathbb{R}[t, \underline{x}, \underline{u}]$ and $\mathbf{X}$ as well as $\mathbf{U}$ are basic, semi-algebraic set, a dual to Putinar's theorem (see [12, Th. 3.8b]) allows to handle measures by their moment sequences. One obtains a finite-dimensional convex relaxation of (3) by truncating the problem with only the first few moments as decision variables, as well as only a few of the linear constraints. Denoting by J^r_{SDP} the cost of the semi-definite relaxation considering moments of degree up to $2r$, [11] have proven

[1] Note that the end-point constraints considered here are for ease of exposition, see [11] for the general case.

that the cost of the relaxations converge from below to the cost of measure LP (3):

Theorem 1 (Lasserre et al.):

$$\lim_{r\to\infty} J^r_{SDP} \uparrow J_{LP}. \tag{4}$$

The following corollary is immediate, and provides the basis for the termination criterion of §III:

Corollary 1: Let an admissible pair $(u(t), x(t))$ have the same cost as a moment relaxation. Then the solution is globally optimal.

For the rest of the paper, we make the following standing assumptions, to ease exposition:

Assumption 1: Problem (1) admits a unique optimal process.

Note that assumption 1 is generic, and can be enforced almost surely by perturbing randomly the coefficients of the polynomials data, as done for polynomial optimization in the `SparsePOP` toolbox [19]. Note that as an immediate corollary, Assumption 1 imposes the uniqueness of the optimal measure by [18, Cor. 1.4], and since measures are moment determinate on compact sets, uniqueness of optimal moment sequences are guaranteed as well. The next sections explore how this unique optimal measure can be reconstructed from its moments data. That is, given definition (2), we propose an algorithm for reconstructing the approximate support of μ.

II. Atomic approximation of occupation measures

This section outlines convergence results for approximating occupation measures by atomic measures, that is, measures supported on a finite number of points only. In the next sections, it is shown how trajectories can be recovered from these approximate measures.

At any given relaxation, only a finite subset of the moment constraints of (3) can be taken into account numerically. As a result, if the optimized moment sequence possesses a representing measure, there always exists an atomic measure with the exact same moments [12, Th. B.12]:

Theorem 2 (Tchakaloff): Let μ be a finite, positive, Borel measure with compact support $\mathbf{Z} \subset \mathbb{R}^q$, and let $d \geq 1$ be a fixed positive integer. Then there exists $p \leq \binom{q+d}{q}$ points $z_k \subset \mathbf{Z}$ and positive weights w_k such that

$$\langle f, \mu\rangle = \sum_{k=1}^{p} w_k\, f(z_k) \tag{5}$$

for every polynomial $f \in \mathbb{R}[z]$ of degree at most d.

To obtain a finite dimensional LP, we therefore propose fixing a time-space grid of Dirac measures, as proposed in [16] for solving (3) directly. As such the only decision variables are then the mass of the Dirac measures, which enter linearly in the problem. The crucial difference here with [16] is that we do not attempt to find an optimal measure, since its (approximate) moments are already given by the semidefinite relaxations. Instead, by segregating optimization and approximation, we circumvent the problem encountered in [16], where there might not be a given atomic measure on a given grid satisfying the truncated moment constraints within a given tolerance. We also prevent the distortion of solutions when prescribed moment tolerances are too loose.

Denote by $\mathbf{Z}_\varepsilon \subset \mathbf{Z} \subset \mathbf{T}\times\mathbf{U}\times\mathbf{X}$ a mesh of given resolution ε supported on finitely many points, such that for all $z \in \mathbf{Z}$, there exists $z_i \in \mathbf{Z}_\varepsilon$ such that $|z_i - z| \leq \varepsilon$. By compactness of $\mathbf{Z}$, such a $\mathbf{Z}_\varepsilon$ always exists.

Theorem 3: Consider the following LP:

$$\begin{aligned} \lambda^*_\varepsilon = \min_{\tilde\mu,\lambda}\ & \lambda \\ \text{s.t. } & |y_\alpha - \langle z^\alpha, \tilde\mu\rangle| \leq \lambda, \quad \forall \alpha \in \mathbb{N}^{1+m+n}_{2r} \\ & \tilde\mu \in \mathcal{M}^+(\mathbf{Z}_\varepsilon), \end{aligned} \tag{6}$$

where y_α are given moments of a representing measure $\mu \in \mathcal{M}^+(\mathbf{Z})$. Then, as the mesh is refined, moments of the approximate atomic measure $\tilde\mu$ converge to those of μ:

$$\lim_{\varepsilon\to 0} \lambda^*_\varepsilon \to 0. \tag{7}$$

Proof: Indeed, Th. 2 asserts the existence of an atomic measure supported on $\binom{1+n+m+2r}{2r}$ points $\{z_k\}$ such that $y_\alpha = \sum_k z_k^\alpha$, for all $\alpha \in \mathbb{N}^{1+n+m}_{2r}$. Therefore, for any $\lambda > 0$, by continuity of the monomials and the finite number of moment constraints, one can always find an $\epsilon > 0$ such that there exists a measure $\tilde\mu \in \mathcal{M}^+(\mathbf{Z}_\epsilon)$ that is admissible for LP (6). As λ is arbitrary and 0 is a lower bound for the cost, this concludes the proof. ∎

See also the related developments in [16, §3], which however try to approximate (3) directly by a LP. To optimize over the cost, moment tolerance λ must then be fixed a priori, and carefully chosen to minimize distortions while guaranteeing the existence of an atomic approximation on the given grid – a very difficult task.

Problem (6) can be interpreted as a truncation of a minimum norm problem, when the weak-$*$ topology is considered for measures. Indeed, on compact sets, it is sufficient to consider a dense basis of the continuous function equipped with the supremum norm. The monomials z^α are such a possible choice, such that the minimum norm problem would read:

$$\inf_{\tilde\mu\in\mathcal{M}^+(\mathbf{Z}_\varepsilon)} \sup_{\alpha\in\mathbb{N}^{1+n+m}} |\langle z^\alpha, \tilde\mu - \mu\rangle|. \tag{8}$$

Clearly, (6) is just the truncation of (8) down to a finite number of moments. This has an immediate consequence by the mere definition of weak-$*$ convergence, as relaxation order r in (6) is increased:

Corollary 2: $\tilde\mu$ converges weakly-$*$ to μ as $r \to \infty$ and $\varepsilon \to 0$.

III. Practical inverse problem

Corollary 2 shows the good asymptotic properties of the algorithm as the grid is refined and the number of moments is increased. Obviously, only a finite number of moments are supplied by any given moment relaxation, and the practical resolution of the LP imposes a maximum size on the approximation grid. From now on, we take a practical view on the inverse problem, working with a finite data set, and with possibly approximate/relaxed moment data as problem input. In the next section, this standing assumption

of "good enough" moments and grid resolution will be tested a posteriori by hot-starting a direct method.

Fix a grid $\mathbf{Z}_\varepsilon$ indexed by q points. Then (6) is rewritten as the following minimum norm problem parametrized by the (positive) weights w_β associated to each atom $z_\beta \in \mathbf{Z}_\varepsilon, \beta = 1, \ldots, q$:

$$\lambda_\epsilon^* = \min_{w \in \mathbb{R}_+^q} \|b - A\,w\|_\infty, \tag{9}$$

with vector b the truncated moment vector $\{y_\alpha\}$, $\alpha \in \mathbb{N}_{2r}^{1+m+n}$ returned by the semi-definite relaxation of order r, and A a $\mathbb{N}_{2r}^{1+m+n} \times q$ matrix, whose elements are given by $A_{\alpha\beta} = z_\beta^\alpha$. This problem can be solved by standard linear programming routines. The support of the atomic measure – which is expected to approximate closely the support of occupation measure (2) – is recovered by the elements of $\mathbf{Z}_\varepsilon$ whose optimal weights are non-zero. By Ass. 1, those waypoints closely approximate the unique optimal trajectory and controls. Note that for the numeric experiments below, it was observed to be more efficient to solve (9) by an interior point method and apply a threshold than to use a simplex-based algorithm.

There is nonetheless a severe restriction to this approach: the dimension of $\mathbf{T}\times\mathbf{X}\times\mathbf{U} \subset \mathbb{R}^{1+n+m}$ leads quickly to LPs of extremely large sizes, a direct consequence of the "curse of dimensionality" of dynamic programming, the conic dual of (3). Indeed, in the opinion of the author, this is the second main restriction for using the numeric approach of [16] for solving (3) directly. Even with a coarse grid of 100 points in each direction, a basic implementation of the algorithm would lead to a LP of a minimum of 10 billion variables for a problem with 1 control and 3 states. Therefore, we propose a middle ground between the rigor and the better scalabilty of the moment approach, and the practical information on the support of the occupation measure given by the LP approach. From the data given by a relaxation, we only consider moments of the form $y_{i0\ldots0j0\ldots0}$. That is we, use only moments of time and one of the states or the controls, such as to perform a coordinate-by-coordinate identification of trajectory and control time series. This way of proceeding keeps the size of the LPs small enough, such that the inverse problem has a negligible resolution time in comparison to that of the moment relaxation, while maintaining a sufficient resolution ε of the grid.

Obviously, although likely to give excellent results even in the face of approximated and truncated data, this procedure does not guarantee admissibility of the reconstructed states and controls for OCP (1). To refine the numeric solutions, we propose to use those approximate time-series to hot-start a direct optimal control method. For all the numeric tests of the next section, we used the freely available software `BOCOP` [1]. Corollary 1 then provides a numeric termination criterion for the moment method: if the cost of the solution given by the direct method is within a prescribed tolerance of the cost of the moment relaxation, the local solution is validated as globally optimal and the hierarchy of moment relaxations can be terminated. However, as shown in the next section, even non-tight relaxations are already enough to recover the global optimal solution via this approach, although those solutions can obviously not be certified to be optimal.

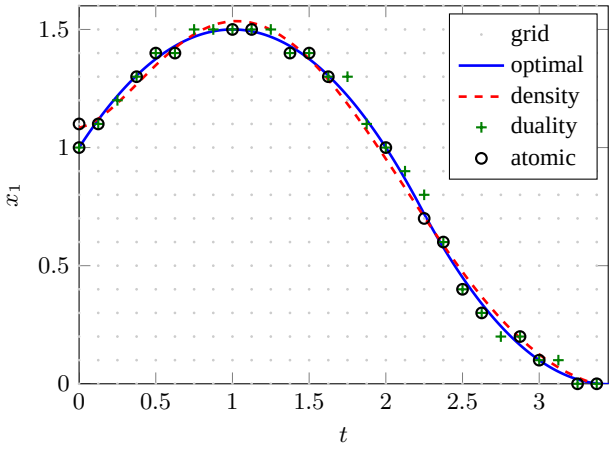

Fig. 2. Reconstruction of state x_1 of the example of §IV-A. Optimal solution versus reconstructions via polynomial densities , dual HJB exploitation and atomic approximations.

IV. Illustrative examples

We present now three illustrative examples. Moment relaxations were solved via `GloptiPoly` [8], using `SeDuMi` [17] as the semi-definite solver.

A. Double integrator

Consider the problem of driving the double integrator to the origin in minimum time:

$$\begin{aligned} J = \inf_{u(t)}\ & t_f \\ \text{s.t. } & \begin{bmatrix}\dot{x}_1\\ \dot{x}_2\end{bmatrix} = \begin{bmatrix}0\\ u\end{bmatrix}, \\ & x(0) = \begin{bmatrix}1\\ 1\end{bmatrix},\ x(t_f) = \begin{bmatrix}0\\ 0\end{bmatrix}, \\ & u(t) \in [-1, 1]. \end{aligned} \tag{10}$$

Figures 2 to 4 present the optimal solution, along with the reconstructed trajectory following a coordinate-by-coordinate identification as proposed in §III. Moments of the fourth-order relaxation were used, since they closely approximate the true moments of the optimal occupation measure. As one can see, the reconstructed trajectories and controls closely match the true solution.

Figures 2 to 4 also compare the proposed atomic reconstruction approach with methods based on the results of [11] and [9]. The approach hinted in [11] consists in extracting dual solutions to each moment relaxation, which give a polynomial (sub)approximation V to the value function satisfying, on $\mathbf{T}\times\mathbf{X}\times\mathbf{U}$,

$$\frac{\partial V}{\partial t} + \frac{\partial V}{\partial x} f + h \geq 0. \tag{11}$$

Around the optimal trajectories, equality in (11) must hold. We therefore propose a very coarse exploitation of such results: for each time t_i in a given discretization of the

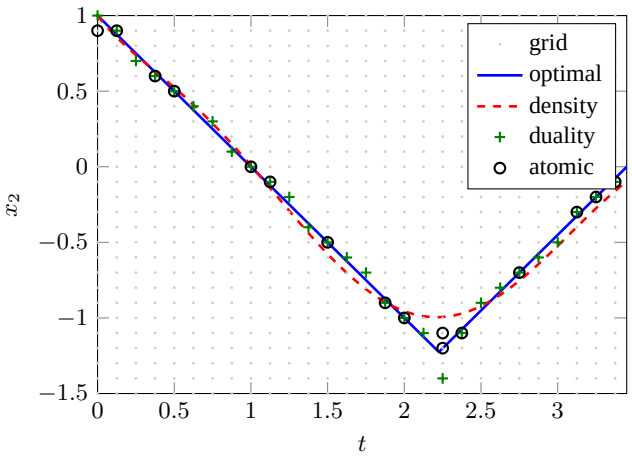

Fig. 3. Reconstruction of state x_2 of the example of §IV-A. Optimal solution versus reconstructions via polynomial densities , dual HJB exploitation and atomic approximations.

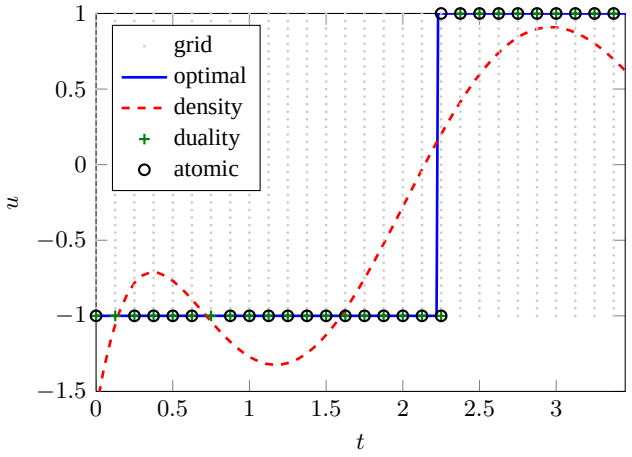

Fig. 4. Reconstruction of control u of the example of §IV-A. Optimal solution versus reconstructions via polynomial densities , dual HJB exploitation and atomic approximations.

time interval, find the couple $(u(t_i), x(t_i))$ minimizing $\frac{\partial V}{\partial t} + \frac{\partial V}{\partial x} f + h$. We did so by exhaustive search on the same time-state-control grid used by the atomic approximation technique, for comparison purposes. Obviously, the main drawback of the approach is the curse of dimensionality, the example presented with 1 control and 2 states being already computationally intensive. The method gives excellent result for reconstructing the state, although slightly less precise than the faster atomic approximation.

Finally, the approach in [9] considers moments of the form $y_{i0\ldots010\ldots0}$ to approximate occupation measures by absolutely continuous (with respect to the Lebesgue measure on time) measures with polynomial densities. This results in a simple linear system to solve. As can be seen in Fig. 2, the method allows for a correct approximation of the first state, while the second state in Fig. 3 cannot capture properly the switch of velocity. The picture is worse for the control, as seen in Fig. 4, as the polynomial density fails to identify any structure such as bang arcs, and gives controls below -1, hence non-admissible. Such behavior is of course expected from a polynomial approximation of a discontinuous function. In addition, this method simply fails if several solutions are encoded in an occupation measures, since the recovered density will be a weighted mean of those solutions. On the other hand, the atomic approximation measure will be supported on all these trajectories. As long as their number is finite, they could be extracted form the approximate measures, albeit with a more sophisticated technique to identify optimal arcs individually.

As a summary, the atomic approximation combines computational effectiveness with a good approximation quality, especially considering admissibility of reconstructed trajectories and detection of control structures.

B. Non-convex integrator

Consider the problem:

$$\begin{aligned} J = \inf_{u(t)} \int_0^1 x^2 \,\mathrm{d}t \\ \text{s.t. } \dot{x} = u, \\ x(0) = 0, \quad x(1) = 3/4, \\ x(t) \in [-1, 1], \\ u(t) \in [1, 1], \\ (x(t) - 1/5)^2 + (t - 1/2)^2 \geq (1/5)^2 . \end{aligned} \tag{12}$$

The last constraint represents a time-dependent obstacle to be avoided, with the cost favoring solution passing below it. However, these solutions cannot be admissible, as the end-point constraints will not be satisfied, and the optimal solution necessarily passes above the obstacle (see Fig. 5). This interplay between non-convex constraints and the cost make this problem very difficult to initialize for local methods, as a few tries with solver BOCOP [1] may reveal.

We now show that the method proposed can be used efficiently to initialize the local method. Figure 5 presents the extracted trajectories from moments of the second relaxation (cost of 0.141), the sixth relaxation (cost of 0.164) and the optimal solution (cost of 0.176). As expected, the sixth relaxation offers a finer reconstruction of the optimal trajectory. However, it should be noted that the very coarse trajectory reconstructed from the second relaxation is already enough to initialize the direct method implemented in BOCOP so that it converges to the global optimum. This suggests that the method proposed in this paper could be widely used for hot starting local methods, as lower-order relaxations are very cheap to compute. This is in contrast with using Hamilton-Jacobi-Bellman solvers as in [4], which offer little possibility of trading solution accuracy with computational load.

C. Van der Pol oscillator

In this last example, we show some alternative use of the proposed approach beyond optimal control. Consider the (uncontrolled) dynamical system $\dot{x} = f(x)$ given by the Van der Pol oscillator:

$$\begin{aligned} \dot{x}_1 &= x_2 \\ \dot{x}_2 &= -x_1 + (1 - x_1^2) x_2 . \end{aligned} \tag{13}$$

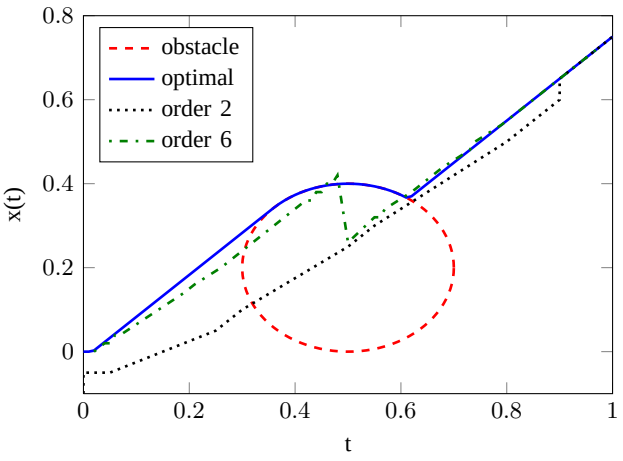

Fig. 5. Comparison between reconstructed trajectories from relaxations of order 2 and 6, against the optimal solution for the example of §IV-B. The latter can be found by using any of the former as a starting point of a local method.

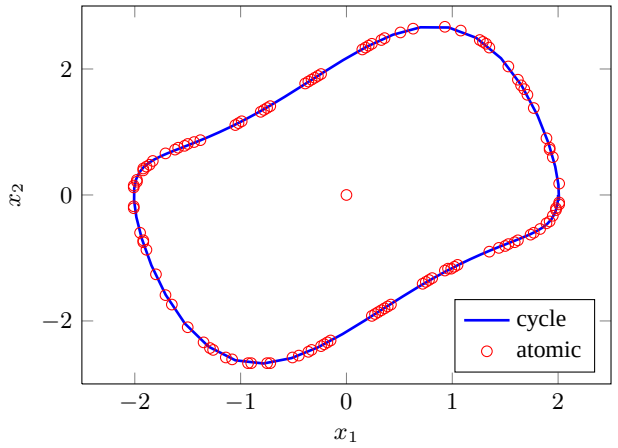

Fig. 6. Reconstruction for the example of §IV-C. Simulated limit cycle versus atomic approximation.

Suppose one would like to compute its limit set, which is known here to be the union of a stable limit cycle and an unstable fixed point. Following e.g. [5], one can relax this as the measure feasibility test

$$\begin{aligned} &\exists \mu ? \\ &\text{s.t. } \forall v \in \mathbb{R}[\underline{x}] : \langle \frac{\partial v}{\partial \underline{x}} f, \mu \rangle = 0, \\ &\quad \langle 1, \mu \rangle = 1, \\ &\quad \mu \in \mathcal{M}^+(\mathbf{X}), \end{aligned} \tag{14}$$

with $\mathbf{X}$ a given compact, basic, semi-algebraic set. This problem can be solved via moment relaxations in a similar procedure as the one presented in §I. Applying the atomic extraction procedure for the support of the resulting measure on $\mathbf{X} = [-3, 3]^2$, and using moments of the 8th relaxation, one obtains the results of Fig. 6. As expected, the extraction procedure locates (within precision of the supplied grid) the stable limit cycle, as well as the unstable equilibrium.

V. Perspectives

In this paper, we present a method for recovering trajectories from the moments of occupation measures. This allows for a numerical method solving optimal control problem globally, in a "black box" fashion. First, moment relaxations as presented in [11] are computed to obtain approximate truncated moment sequences. Then, the method outlined here solves the inverse problem to obtain approximate global solutions. Those solutions are then fed to a local method to guarantee admissibility of the process.

Numerically, the inverse problem (9) possesses a simple structure, as one tries to minimize the distance in a finite-dimensional space, as measured by the supremum norm, from a point to a linear subspace. As such, it is expected that one could use dedicated solvers, using large scale, problem-specific routines to outperform the simple linear programming method considered here. See e.g. [13] for an introduction on the subject. This opens the possibility of identifying approximate trajectories and controls in one pass, using all information available from moments.

Finally, we also presented the approach for bounded control for ease of exposition, but the method holds as well for modal occupation measures of switched systems [6] and impulsive occupation measures [3]. This demonstrates the high flexibility of the approach.

Acknowledgments

This work was supported by the Engineering and Physical Sciences Research Council under Grant EP/G066477/1, and benefited from discussions with Jean-Bernard Lasserre and Pierre Martinon.

References

[1] F. J. Bonnans, P. Martinon, V. Grélard. Bocop - A collection of examples. Technical report RR-8053, INRIA, 2012.
[2] A. Cernea, H. Frankowska. A Connection Between the Maximum Principle and Dynamic Programming for Constrained Control Problems. SIAM J. Control Optim., 44(2):673–703, 2005.
[3] M. Claeys, D. Arzelier, D. Henrion, J. B. Lasserre. Measures and LMIs for Impulsive Nonlinear Optimal Control. IEEE Trans. Autom. Control, 2014. Accepted for publication.
[4] E. Cristiani, P. Martinon. Initialization of the shooting method via the Hamilton-Jacobi-Bellman approach. J. of Optim. Theory and Applications, 146:321–346, 2010.
[5] V. Gaitsgory, M. Quincampoix. Linear programming approach to deterministic infinite horizon optimal control problems with discounting. SIAM J. Control Optim., 48(4): 2480–2512, 2009.
[6] D. Henrion, J. Daafouz, M. Claeys. Optimal switching control design for polynomial systems: an LMI approach. Proc. Conf. Decision Control, Florence, Italy, 2013.
[7] D. Henrion, M. Korda. Convex computation of the region of attraction of polynomial control systems. LAAS-CNRS Research Report 12488, 2012. To appear in the IEEE Trans. Automatic Control.
[8] D. Henrion, J. B. Lasserre, J. Löfberg. GloptiPoly 3: Moments, Optimization and Semidefinite Programming. Optim. Methods and Software, 24(4-5):761–779, 2009.
[9] D. Henrion, J. B. Lasserre, M. Mevissen. Mean squared error minimization for inverse moment problems. LAAS-CNRS Research Report 12487, 2012. Submitted to Applied Mathematics and Optim., 2012.
[10] J. B. Lasserre. Optimisation globale et théorie des moments, Comptes Rendus de l'Académie des Sciences - Series I - Mathematics, 331(11), 929–934, 2000.
[11] J. B. Lasserre, D. Henrion, C. Prieur, E. Trélat. Nonlinear optimal control via occupation measures and LMI relaxations. SIAM J. Control Opt., 47:1643-1666, 2008.

[12] J. B. Lasserre. Moments, positive polynomials and their applications. Imperial College Press, London, UK, 2009.
[13] Y. Nesterov. Random gradient-free minimization of convex functions. CORE discussion paper 2011/1, Université Catholique de Louvain , 2011.
[14] J. Nie. Optimality Conditions and Finite Convergence of Lasserre's Hierarchy. Mathematical Programming, Ser. A, 2013.
[15] L. S. Pontryagin, V. G. Boltyanskii, R. V. Gamkrelidze, E. F. Mishchenko. The mathematical theory of optimal processes. Interscience, 1962.
[16] J. E. Rubio. Control and optimization: the linear treatment of nonlinear problems. Manchester University Press, UK, 1986.
[17] J.F. Sturm. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. Optim. Methods and Software, 11-12:625–653, 1999.
[18] R. Vinter. Convex Duality and Nonlinear Optimal Control. SIAM J. Control Optim., 31:518-538, 1993.
[19] H. Waki, S. Kim, M. Kojima, M. Muramatsu and H. Sugimoto, SparsePOP : a Sparse Semidefinite Programming Relaxation of Polynomial Optimization Problems, ACM Trans. Mathematical Software , 35(2):15, 2008.